Remarks and Arguments

Claims 1-11 and 44-54 were presented for examination. Claims 1 and 44 have been amended.

Applicant's attorney thanks the examiner for his time during a telephone interview on August 16, 2007. During that interview the application of the Lumelsky and Schaefer references to claim 1 was discussed as well as the operation of inventive system in dynamically varying the locations in which data is stored based on criteria such as workload and availability. These discussions are set out in detail below.

Claims 1-11 and 44-54 have been rejected under 35 U.S.C. §103(a) as obvious over U.S. Patent No. 6,460,082 (Lumelsky) in view of U.S. Patent No. 6,157,927 (Schaefer.) As discussed in the telephone interview, in conventional data storage systems that were available at the time that the application was filed, an address which identified the data directly corresponded to the actual physical location at which the data was to be stored. Therefore, during a typical read-modify-write sequence, data identified by an address was retrieved from a location specified by that address, modified and then stored back to the same location as specified by that address. This type of operation had several consequences. In particular, it made it difficult to effectively load-balance data storage systems that consisted of multiple discrete storage units such as disks. A typical scenario was a data warehousing application in which a large number of disks were used to service the storage needs of a plurality of customers. It was the typical practice in such systems to allocate a number of disks to each customer based on criteria such as guaranteed quality of service. This allowed a customer's data to be retrieved and stored in the same locations as discussed above. However, it also meant that the storage resources devoted to one customer might be overloaded while the storages resource of another customer sat idle thereby violating the quality of service guarantees. This was particularly true when a disk failed or was taken out of service for maintenance, etc. Therefore, such a system required constant manual shifting of storage resources between customers to maintain the quality of service guarantees.

The present invention alleviated this problem by breaking the correspondence between data identification and the location at which the data is stored. In particular, in

the present system, a client wishing to retrieve or store data specifies a data identifier instead of a traditional address that is associated with a location. This data identifier is presented to an access interface which determines a location at which the data will be stored, based not on the identity of the data, but on other criteria such as the workload of each resource in the resource pool. The result of this operation is that data may be stored in a storage resource other than the storage resource from which that data was retrieved. Therefore, during a typical read-modify-write operation, data corresponding a data identifier would be read from a storage location and modified, but might be stored back into a location different from which it was read depending on resource load and availability. The data thus dynamically moves based on criteria determined by the data storage system. Consequently, load balancing becomes automatic.

The Lumelsky reference discloses a distributed multimedia system in which clients receive streaming multimedia information from a distributed pool of multimedia servers. In this system, the same data may be retrieved from different resources, based on criteria such as load balancing. However, the direction of data flow is always from server to client. No data is received from clients and transferred to the multimedia servers and stored thereon – no storage write operations are performed.

As noted in the telephone interview, the examiner interpreted the Schaefer reference as disclosing two units connected by a switching network and commented that at least one of these units, the transaction manager, could correspond to the access interface recited in the claims in that it at least temporarily stored data before transferring that data over the switching system to a storage resource. While applicants agree that Schaefer discloses two units connected by a switching network, as previously noted, nothing in Schaefer indicates that the system uses anything other than conventional database transaction processing in which information is retrieved from, and stored back into, the same database location. Specifically, the storage location is determined by the address and not by the workload of the storage resources.

Therefore, neither of the cited references discloses a system in which data corresponding to the same data identifier may be stored in different storage locations based on criteria, such as load balancing or availability and, consequently, the combination of these references cannot teach or suggest this type of operation. The

claims have been amended to particularly point out the differences between the present system and the cited references. For example, claim 1 now recites, in lines 3-10, "…an access interface module which receives data storage requests, each including a data identifier, and data to be stored from the client and, in response to each service request and based on a workload instead of a location in the plurality of resources, dynamically selects a subset of the plurality of storage resources to which the data is transferred to be stored so that the location to which data corresponding to the same data identifier is transferred can change from request to request in order to dynamically distribute the workload across the plurality of storage resources…" (emphasis added). As discussed above, in the Lumelsky system, no data storage operations are disclosed. In the Schaefer system, the resources to which data is transferred are selected by address not by workload as claimed. Thus, the combination of Lumelsky and Schaefer references proposed by the examiner cannot teach or suggest the structure recited in claim 1 because neither references teaches this structure. Therefore, claim 1 patentably distinguishes over the cited reference combination.

Claims 2-11 are dependent, either directly or indirectly, on amended claim 1 and incorporate the limitations thereof. Therefore, they also distinguish over the cited reference combination in the same manner as amended claim 1.

Method claim 44 contains limitations that parallel those in apparatus claim 1 and has been amended in a manner that parallels that of claim 1. Thus, amended claim 44 patentably distinguishes over the cited reference combination in the same manner as claim 1. Claims 45-54 are dependent, either directly or indirectly, on amended claim 44 and incorporate the limitations thereof. Therefore, they also distinguish over the cited reference combination in the same manner as amended claim 44.

In light of the forgoing amendments and remarks, this application is now believed in condition for allowance and a notice of allowance is earnestly solicited. If the examiner has any further questions regarding this amendment, he is invited to call applicants' attorney at the number listed below. The examiner is hereby authorized to charge any fees or direct any payment under 37 C.F.R. §§1.17, 1.16 to Deposit Account number 50-3969.

Respectfully submitted

/paul e. kudirka/ Date: 2007-09-26

Paul E. Kudirka, Esq. Reg. No. 26,931
LAW OFFICES OF PAUL E. KUDIRKA
Customer Number 64967
Tel: (617) 357-0010 Fax: (617) 357-0035